

Elusive knowledge

We know a lot. I know what food penguins eat. I know that phones used to ring, but nowadays squeal, when someone calls up. I know that Essendon won the 1993 Grand Final. I know that here is a hand, and here is another.

We have all sorts of everyday knowledge, and we have it in abundance. To doubt that would be absurd. At any rate, to doubt it in any serious and lasting way would be absurd; and even philosophical and temporary doubt, under the influence of argument, is more than a little peculiar. It is a Moorean fact that we know a lot. It is one of those things that we know better than we know the premises of any philosophical argument to the contrary.

Besides knowing a lot that is everyday and trite, I myself think that we know a lot that is interesting and esoteric and controversial. We know a lot about things unseen: tiny particles and pervasive fields, not to mention one another's underwear. Sometimes we even know what an author meant by his writings. But on these questions, let us agree to disagree peacefully with the champions of 'post-knowledgeism'.

First published in *The Australasian Journal of Philosophy* 74 (1996), pp. 549–567. Reprinted with kind permission from *The Australasian Journal of Philosophy*.

Thanks to many for valuable discussions of this material. Thanks above all to Peter Unger; and to Stewart Cohen, Michael Devitt, Alan Hájek, Stephen Hetherington, Denis Robinson, Ernest Sosa, Robert Stalnaker, Jonathan Vogel, and a referee for *The Australasian Journal of Philosophy*. Thanks also to the Boyce Gibson Memorial Library and to Ormond College.

The most trite and ordinary parts of our knowledge will be problem enough.

For no sooner do we engage in epistemology – the systematic philosophical examination of knowledge – than we meet a compelling argument that we know next to nothing. The sceptical argument is nothing new or fancy. It is just this: it seems as if knowledge must be by definition infallible. If you claim that *S* knows that *P*, and yet you grant that *S* cannot eliminate a certain possibility in which not-*P*, it certainly seems as if you have granted that *S* does not after all know that *P*. To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* contradictory.

Blind Freddy can see where this will lead. Let your paranoid fantasies rip – CIA plots, hallucinogens in the tap water, conspiracies to deceive, old Nick himself – and soon you find that uneliminated possibilities of error are everywhere. Those possibilities of error are far-fetched, of course, but possibilities all the same. They bite into even our most everyday knowledge. We never have infallible knowledge.

Never – well, hardly ever. Some say we have infallible knowledge of a few simple, axiomatic necessary truths; and of our own present experience. They say that I simply cannot be wrong that a part of a part of something is itself a part of that thing; or that it seems to me now (as I sit here at the keyboard) exactly as if I am hearing clicking noises on top of a steady whirring. Some say so. Others deny it. No matter; let it be granted, at least for the sake of the argument. It is not nearly enough. If we have only that much infallible knowledge, yet knowledge is by definition infallible, then we have very little knowledge indeed – not the abundant everyday knowledge we thought we had. That is still absurd.

So we know a lot; knowledge must be infallible; yet we have fallible knowledge or none (or next to none). We are caught between the rock of fallibilism and the whirlpool of scepticism. Both are mad!

Yet fallibilism is the less intrusive madness. It demands less frequent corrections of what we want to say. So, if forced to choose, I choose fallibilism. (And so say all of us.) We can get used to it, and some of us have done. No joy there – we know that people can get used to the most crazy philosophical sayings imaginable. If you are a contented

fallibilist, I implore you to be honest, be naive, hear it afresh. ‘He knows, yet he has not eliminated all possibilities of error.’ Even if you’ve numbed your ears, doesn’t this overt, explicit fallibilism *still* sound wrong?

Better fallibilism than scepticism; but it would be better still to dodge the choice. I think we can. We will be alarmingly close to the rock, and also alarmingly close to the whirlpool, but if we steer with care, we can – just barely – escape them both.

Maybe epistemology is the culprit. Maybe this extraordinary pastime robs us of our knowledge. Maybe we do know a lot in daily life; but maybe when we look hard at our knowledge, it goes away. But only when we look at it harder than the sane ever do in daily life; only when we let our paranoid fantasies rip. That is when we are forced to admit that there always are uneliminated possibilities of error, so that we have fallible knowledge or none.

Much that we say is context-dependent, in simple ways or subtle ways. Simple: ‘it’s evening’ is truly said when, and only when, it is said in the evening. Subtle: it could well be true, and not just by luck, that Essendon played rottenly, the Easybeats played brilliantly, yet Essendon won. Different contexts evoke different standards of evaluation. Talking about the Easybeats we apply lax standards, else we could scarcely distinguish their better days from their worse ones. In talking about Essendon, no such laxity is required. Essendon won because play that is rotten by demanding standards suffices to beat play that is brilliant by lax standards.

Maybe ascriptions of knowledge are subtly context-dependent, and maybe epistemology is a context that makes them go false. Then epistemology would be an investigation that destroys its own subject matter. If so, the sceptical argument might be flawless, when we engage in epistemology – and only then!¹

1 The suggestion that ascriptions of knowledge go false in the context of epistemology is to be found in Barry Stroud, ‘Understanding Human Knowledge in General’ in Marjorie Clay and Keith Lehrer (eds.), *Knowledge and Skepticism* (Boulder: Westview Press, 1989); and in Stephen Hetherington, ‘Lacking Knowledge and Justification by Theorising About Them’ (lecture at the University of New South Wales, August

If you start from the ancient idea that justification is the mark that distinguishes knowledge from mere opinion (even true opinion), then you well might conclude that ascriptions of knowledge are context-dependent because standards for adequate justification are context-dependent. As follows: opinion, even if true, deserves the name of knowledge only if it is adequately supported by reasons; to deserve that name in the especially demanding context of epistemology, the arguments from supporting reasons must be especially watertight; but the special standards of justification that this special context demands never can be met (well, hardly ever). In the strict context of epistemology we know nothing, yet in laxer contexts we know a lot.

But I myself cannot subscribe to this account of the context-dependence of knowledge, because I question its starting point. I don't agree that the mark of knowledge is justification.² First, because justification is not sufficient: your true opinion that you will lose the lottery isn't knowledge, whatever the odds. Suppose you know that it is a fair lottery with one winning ticket and many losing tickets, and you know how many losing tickets there are. The greater the number of losing tickets, the better is your justification for believing you will lose. Yet there is no number great enough to transform your fallible opinion into knowledge – after all, you just might win. No justification is good enough – or none short of a watertight deductive argument, and all but the sceptics will agree that this is too much to demand.³

Second, because justification is not always necessary. What (non-circular) argument supports our reliance on perception, on memory,

1992). Neither of them tells the story just as I do, however it may be that their versions do not conflict with mine.

- 2 Unless, like some, we simply define 'justification' as 'whatever it takes to turn true opinion into knowledge' regardless of whether what it takes turns out to involve argument from supporting reasons.
- 3 The problem of the lottery was introduced in Henry Kyburg, *Probability and the Logic of Rational Belief* (Middletown, CT: Wesleyan University Press, 1961), and in Carl Hempel, 'Deductive-Nomological vs. Statistical Explanation' in Herbert Feigl and Grover Maxwell (eds.), *Minnesota Studies in the Philosophy of Science*, Vol. II (Minneapolis: University of Minnesota Press, 1962). It has been much discussed since, as a problem both about knowledge and about our everyday, non-quantitative concept of belief.

and on testimony?⁴ And yet we do gain knowledge by these means. And sometimes, far from having supporting arguments, we don't even know how we know. We once had evidence, drew conclusions, and thereby gained knowledge; now we have forgotten our reasons, yet still we retain our knowledge. Or we know the name that goes with the face, or the sex of the chicken, by relying on subtle visual cues, without knowing what those cues may be.

The link between knowledge and justification must be broken. But if we break that link, then it is not – or not entirely, or not exactly – by raising the standards of justification that epistemology destroys knowledge. I need some different story.

To that end, I propose to take the infallibility of knowledge as my starting point.⁵ Must infallibilist epistemology end in scepticism? Not quite. Wait and see. Anyway, here is the definition. Subject *S* *knows* proposition *P* iff *P* holds in every possibility left uneliminated by *S*'s evidence; equivalently, iff *S*'s evidence eliminates every possibility in which not-*P*.

The definition is short, the commentary upon it is longer. In the first place, there is the proposition, *P*. What I choose to call 'propositions' are individuated coarsely, by necessary equivalence. For instance, there is only one necessary proposition. It holds in every possibility; hence in every possibility left uneliminated by *S*'s evidence, no matter who *S* may be and no matter what his evidence may be. So the necessary proposition is known always and everywhere. Yet this known proposition may go unrecognised when presented in impenetrable linguistic disguise, say as the proposition that every even number is the sum of two primes. Likewise, the known proposition that I have two hands may go unrecognised when presented as the proposition that the number of my hands is the least number *n* such that every even number is the sum of *n* primes. (Or if you doubt the necessary existence of numbers, switch to an example involving equiv-

4 The case of testimony is less discussed than the others; but see C. A. J. Coady, *Testimony: A Philosophical Study* (Oxford: Clarendon Press, 1992) pp. 79–129.

5 I follow Peter Unger, *Ignorance: A Case for Skepticism* (New York: Oxford University Press, 1975). But I shall not let him lead me into scepticism.

alence by logic alone.) These problems of disguise shall not concern us here. Our topic is modal, not hyperintensional, epistemology.⁶

Next, there are the possibilities. We needn't enter here into the question whether these are *concreta*, abstract constructions, or abstract simples. Further, we needn't decide whether they must always be maximally specific possibilities, or whether they need only be specific enough for the purpose at hand. A possibility will be specific enough if it cannot be split into subcases in such a way that anything we have said about possibilities, or anything we are going to say before we are done, applies to some subcases and not to others. For instance, it should never happen that proposition *P* holds in some but not all sub-cases; or that some but not all sub-cases are eliminated by *S*'s evidence.

But we do need to stipulate that they are not just possibilities as to how the whole world is; they also include possibilities as to which part of the world is oneself, and as to when it now is. We need these possibilities *de se et nunc* because the propositions that may be known include propositions *de se et nunc*.⁷ Not only do I know that there are hands in this world somewhere and somewhen. I know that *I* have hands, or anyway I have them *now*. Such propositions aren't just made true or made false by the whole world once and for all. They are true for some of us and not for others, or true at some times and not others, or both.

Further, we cannot limit ourselves to 'real' possibilities that conform to the actual laws of nature, and maybe also to actual past history. For propositions about laws and history are contingent, and may or may not be known.

Neither can we limit ourselves to 'epistemic' possibilities for *S* – possibilities that *S* does not know not to obtain. That would drain our definition of content. Assume only that knowledge is closed under strict implication. (We shall consider the merits of this assumption later.) Remember that we are not distinguishing between equivalent

6 See Robert Stalnaker, *Inquiry* (Cambridge, MA: MIT Press, 1984) pp. 59–99.

7 See my 'Attitudes *De Dicto* and *De Se*', *The Philosophical Review* 88 (1979) pp. 513–543; and R. M. Chisholm, 'The Indirect Reflexive' in C. Diamond and J. Teichman (eds.), *Intention and Intentionality: Essays in Honour of G. E. M. Anscombe* (Brighton: Harvester, 1979).

propositions. Then knowledge of a conjunction is equivalent to knowledge of every conjunct. P is the conjunction of all propositions not- W , where W is a possibility in which not- P . That suffices to yield an equivalence: S knows that P iff, for every possibility W in which not- P , S knows that not- W . Contrapositing and cancelling a double negation: iff every possibility which S does not know not to obtain is one in which P . For short: iff P holds throughout S 's epistemic possibilities. Yet to get this far, we need no substantive definition of knowledge at all! To turn this into a substantive definition, in fact the very definition we gave before, we need to say one more thing: S 's epistemic possibilities are just those possibilities that are uneliminated by S 's evidence.

So, next, we need to say what it means for a possibility to be eliminated or not. Here I say that the uneliminated possibilities are those in which the subject's entire perceptual experience and memory are just as they actually are. There is one possibility that actually obtains (for the subject and at the time in question); call it *actuality*. Then a possibility W is *uneliminated* iff the subject's perceptual experience and memory in W exactly match his perceptual experience and memory in actuality. (If you want to include other alleged forms of basic evidence, such as the evidence of our extrasensory faculties, or an innate disposition to believe in God, be my guest. If they exist, they should be included. If not, no harm done if we have included them conditionally.)

Note well that we do not need the 'pure sense-datum language' and the 'incorrigible protocol statements' that for so long bedevilled foundationalist epistemology. It matters not at all whether there are words to capture the subject's perceptual and memory evidence, nothing more and nothing less. If there are such words, it matters not at all whether the subject can hit upon them. The given does not consist of basic axioms to serve as premises in subsequent arguments. Rather, it consists of a match between possibilities.

When perceptual experience E (or memory) eliminates a possibility W , that is not because the propositional content of the experience conflicts with W . (Not even if it is the narrow content.) The propositional content of our experience could, after all, be false. Rather, it is the existence of the experience that conflicts with W : W is a possibility in which the subject is not having experience E . Else we would need to

tell some fishy story of how the experience has some sort of infallible, ineffable, purely phenomenal propositional content. . . . Who needs that? Let E have propositional content P . Suppose even – something I take to be an open question – that E is, in some sense, fully characterized by P . Then I say that E eliminates W iff W is a possibility in which the subject's experience or memory has content different from P . I do *not* say that E eliminates W iff W is a possibility in which P is false.

Maybe not every kind of sense perception yields experience; maybe, for instance, the kinaesthetic sense yields not its own distinctive sort of sense-experience but only spontaneous judgements about the position of one's limbs. If this is true, then the thing to say is that kinaesthetic evidence eliminates all possibilities except those that exactly resemble actuality with respect to the subject's spontaneous kinaesthetic judgements. In saying this, we would treat kinaesthetic evidence more on the model of memory than on the model of more typical senses.

Finally, we must attend to the word 'every'. What does it mean to say that every possibility in which not- P is eliminated? An idiom of quantification, like 'every', is normally restricted to some limited domain. If I say that every glass is empty, so it's time for another round, doubtless I and my audience are ignoring most of all the glasses there are in the whole wide world throughout all of time. They are outside the domain. They are irrelevant to the truth of what was said.

Likewise, if I say that every uneliminated possibility is one in which P , or words to that effect, I am doubtless ignoring some of all the uneliminated alternative possibilities that there are. They are outside the domain, they are irrelevant to the truth of what was said.

But, of course, I am not entitled to ignore just any possibility I please. Else true ascriptions of knowledge, whether to myself or to others, would be cheap indeed. I may properly ignore some uneliminated possibilities; I may not properly ignore others. Our definition of knowledge requires a *sotto voce* proviso. S knows that P iff S 's evidence eliminates every possibility in which not- P – Psst! – except for those possibilities that we are properly ignoring.

Unger suggests an instructive parallel.⁸ Just as P is known iff there are no uneliminated possibilities of error, so likewise a surface is flat

⁸ Peter Unger, *Ignorance*, chapter II. I discuss the case, and briefly foreshadow the

iff there are no bumps on it. We must add the proviso: Psst! – except for those bumps that we are properly ignoring. Else we will conclude, absurdly, that nothing is flat. (Simplify by ignoring departures from flatness that consist of gentle curvature.)

We can restate the definition. Say that we *presuppose* proposition Q iff we ignore all possibilities in which not-Q. To close the circle: we *ignore* just those possibilities that falsify our presuppositions. *Proper* presupposition corresponds, of course, to proper ignoring. Then S knows that P iff S's evidence eliminates every possibility in which not-P – Psst! – except for those possibilities that conflict with our proper presuppositions.⁹

The rest of (modal) epistemology examines the *sotto voce* proviso. It asks: what may we properly presuppose in our ascriptions of knowledge? Which of all the uneliminated alternative possibilities may not properly be ignored? Which ones are the ‘relevant alternatives’? – relevant, that is, to what the subject does and doesn’t know?¹⁰ In reply, we can list several rules.¹¹ We begin with three prohibitions: rules to tell us what possibilities we may not properly ignore.

First, there is the *Rule of Actuality*. The possibility that actually obtains is never properly ignored; actuality is always a relevant alternative;

present paper, in my ‘Scorekeeping in a Language Game’, *Journal of Philosophical Logic* 8 (1979) pp. 339–359, esp. pp. 353–355.

9 See Robert Stalnaker, ‘Presuppositions’, *Journal of Philosophical Logic* 2 (1973) pp. 447–457; and ‘Pragmatic Presuppositions’ in Milton Munitz and Peter Unger (eds.), *Semantics and Philosophy* (New York: New York University Press, 1974). See also my ‘Scorekeeping in a Language Game’.

The definition restated in terms of presupposition resembles the treatment of knowledge in Kenneth S. Ferguson, *Philosophical Scepticism* (Cornell University doctoral dissertation, 1980).

10 See Fred Dretske, ‘Epistemic Operators’, *The Journal of Philosophy* 67 (1970) pp. 1007–1022, and ‘The Pragmatic Dimension of Knowledge’, *Philosophical Studies* 40 (1981) pp. 363–378; Alvin Goldman, ‘Discrimination and Perceptual Knowledge’, *The Journal of Philosophy* 73 (1976) pp. 771–791; G. C. Stine, ‘Skepticism, Relevant Alternatives, and Deductive Closure’, *Philosophical Studies* 29 (1976) pp. 249–261; and Stewart Cohen, ‘How to be A Fallibilist’, *Philosophical Perspectives* 2 (1988) pp. 91–123.

11 Some of them, but only some, taken from the authors just cited.

nothing false may properly be presupposed. It follows that only what is true is known, wherefore we did not have to include truth in our definition of knowledge. The rule is ‘externalist’ – the subject himself may not be able to tell what is properly ignored. In judging which of his ignorings are proper, hence what he knows, we judge his success in knowing – not how well he tried.

When the Rule of Actuality tells us that actuality may never be properly ignored, we can ask: *whose* actuality? Ours, when we ascribe knowledge or ignorance to others? Or the subject’s? In simple cases, the question is silly. (In fact, it sounds like the sort of pernicious nonsense we would expect from someone who mixes up what is true with what is believed.) There is just one actual world, we the ascribers live in that world, the subject lives there too, so the subject’s actuality is the same as ours.

But there are other cases, less simple, in which the question makes perfect sense and needs an answer. Someone may or may not know who he is; someone may or may not know what time it is. Therefore I insisted that the propositions that may be known must include propositions *de se et nunc*; and likewise that the possibilities that may be eliminated or ignored must include possibilities *de se et nunc*. Now we have a good sense in which the subject’s actuality may be different from ours. I ask today what Fred knew yesterday. In particular, did he then know who he was? Did he know what day it was? Fred’s actuality is the possibility *de se et nunc* of being Fred on September 19th at such-and-such possible world; whereas my actuality is the possibility *de se et nunc* of being David on September 20th at such-and-such world. So far as the world goes, there is no difference: Fred and I are worldmates, his actual world is the same as mine. But when we build subject and time into the possibilities *de se et nunc*, then his actuality yesterday does indeed differ from mine today.

What is more, we sometimes have occasion to ascribe knowledge to those who are off at other possible worlds. I didn’t read the newspaper yesterday. What would I have known if I had read it? More than I do in fact know. (More and less: I do in fact know that I left the newspaper unread, but if I had read it, I would not have known that I had left it unread.) I-who-did-not-read-the-newspaper am

here at this world, ascribing knowledge and ignorance. The subject to whom I am ascribing that knowledge and ignorance, namely I-as-I-would-have-been-if-I-had-read-the-newspaper, is at a different world. The worlds differ in respect at least of a reading of the newspaper. Thus the ascriber's actual world is not the same as the subject's. (I myself think that the ascriber and the subject are two different people: the subject is the ascriber's otherworldly counterpart. But even if you think the subject and the ascriber are the same identical person, you must still grant that this person's actuality *qua* subject differs from his actuality *qua* ascriber.)

Or suppose we ask modal questions about the subject: what must he have known, what might he have known? Again we are considering the subject as he is not here, but off at other possible worlds. Likewise if we ask questions about knowledge of knowledge: what does he (or what do we) know that he knows?

So the question 'whose actuality?' is not a silly question after all. And when the question matters, as it does in the cases just considered, the right answer is that it is the subject's actuality, not the ascriber's, that never can be properly ignored.

Next, there is the *Rule of Belief*. A possibility that the subject believes to obtain is not properly ignored, whether or not he is right to so believe. Neither is one that he ought to believe to obtain – one that evidence and arguments justify him in believing – whether or not he does so believe.

That is rough. Since belief admits of degree, and since some possibilities are more specific than others, we ought to reformulate the rule in terms of degree of belief, compared to a standard set by the unspecificity of the possibility in question. A possibility may not be properly ignored if the subject gives it, or ought to give it, a degree of belief that is sufficiently high, and high not just because the possibility in question is unspecific.

How high is 'sufficiently high'? That may depend on how much is at stake. When error would be especially disastrous, few possibilities may be properly ignored. Then even quite a low degree of belief may be 'sufficiently high' to bring the Rule of Belief into play. The jurors

know that the accused is guilty only if his guilt has been proved beyond reasonable doubt.¹²

Yet even when the stakes are high, some possibilities still may be properly ignored. Disastrous though it would be to convict an innocent man, still the jurors may properly ignore the possibility that it was the dog, marvellously well-trained, that fired the fatal shot. And, unless they are ignoring other alternatives more relevant than that, they may rightly be said to know that the accused is guilty as charged. Yet if there had been reason to give the dog hypothesis a slightly less negligible degree of belief – if the world's greatest dog-trainer had been the victim's mortal enemy – then the alternative would be relevant after all.

This is the only place where belief and justification enter my story. As already noted, I allow justified true belief without knowledge, as in the case of your belief that you will lose the lottery. I allow knowledge without justification, in the cases of face recognition and chicken sexing. I even allow knowledge without belief, as in the case of the timid student who knows the answer but has no confidence that he has it right, and so does not believe what he knows.¹³ Therefore any proposed converse to the Rule of Belief should be rejected. A possibility that the subject does not believe to a sufficient degree, and ought not to believe to a sufficient degree, may nevertheless be a relevant alternative and not properly ignored.

Next, there is the *Rule of Resemblance*. Suppose one possibility saliently resembles another. Then if one of them may not be properly ignored, neither may the other. (Or rather, we should say that if one of them may not properly be ignored *in virtue of rules other than this rule*, then neither may the other. Else nothing could be properly ignored; because enough little steps of resemblance can take us from anywhere to any-

12 Instead of complicating the Rule of Belief as I have just done, I might equivalently have introduced a separate Rule of High Stakes saying that when error would be especially disastrous, few possibilities are properly ignored.

13 A. D. Woozley, 'Knowing and Not Knowing', *Proceedings of the Aristotelian Society* 53 (1953) pp. 151–172; Colin Radford, 'Knowledge – by Examples', *Analysis* 27 (1966) pp. 1–11.

where.) Or suppose one possibility saliently resembles two or more others, one in one respect and another in another, and suppose that each of these may not properly be ignored (in virtue of rules other than this rule). Then these resemblances may have an additive effect, doing more together than any one of them would separately.

We must apply the Rule of Resemblance with care. Actuality is a possibility uneliminated by the subject's evidence. Any other possibility *W* that is likewise uneliminated by the subject's evidence thereby resembles actuality in one salient respect: namely, in respect of the subject's evidence. That will be so even if *W* is in other respects very dissimilar to actuality – even if, for instance, it is a possibility in which the subject is radically deceived by a demon. Plainly, we dare not apply the Rules of Actuality and Resemblance to conclude that any such *W* is a relevant alternative – that would be capitulation to scepticism. The Rule of Resemblance was never meant to apply to *this* resemblance! We seem to have an *ad hoc* exception to the Rule, though one that makes good sense in view of the function of attributions of knowledge. What would be better, though, would be to find a way to reformulate the Rule so as to get the needed exception without *ad hoc*ery. I do not know how to do this.

It is the Rule of Resemblance that explains why you do not know that you will lose the lottery, no matter what the odds are against you and no matter how sure you should therefore be that you will lose. For every ticket, there is the possibility that it will win. These possibilities are saliently similar to one another: so either every one of them may be properly ignored, or else none may. But one of them may not properly be ignored: the one that actually obtains.

The Rule of Resemblance also is the rule that solves the Gettier problems: other cases of justified true belief that are not knowledge.¹⁴

14 See Edmund Gettier, 'Is Justified True Belief Knowledge?', *Analysis* 23 (1963) pp. 121–123. Diagnoses have varied widely. The four examples below come from: (1) Keith Lehrer and Thomas Paxson Jr., 'Knowledge: Undefeated True Belief', *The Journal of Philosophy* 66 (1969) pp. 225–237; (2) Bertrand Russell, *Human Knowledge: Its Scope and Limits* (London: Allen and Unwin, 1948) p. 154; (3) Alvin Goldman, 'Discrimination and Perceptual Knowledge', op. cit.; (4) Gilbert Harman, *Thought* (Princeton, NJ: Princeton University Press, 1973) p. 143.

Though the lottery problem is another case of justified true belief without

(1) I think that Nogot owns a Ford, because I have seen him driving one; but unbeknownst to me he does not own the Ford he drives, or any other Ford. Unbeknownst to me, Havit does own a Ford, though I have no reason to think so because he never drives it, and in fact I have often seen him taking the tram. My justified true belief is that one of the two owns a Ford. But I do not know it; I am right by accident. Diagnosis: I do not know, because I have not eliminated the possibility that Nogot drives a Ford he does not own whereas Havit neither drives nor owns a car. This possibility may not properly be ignored. Because, first, actuality may not properly be ignored; and, second, this possibility saliently resembles actuality. It resembles actuality perfectly so far as Nogot is concerned; and it resembles actuality well so far as Havit is concerned, since it matches actuality both with respect to Havit's careless habits and with respect to the general correlation between careless habits and carelessness. In addition, this possibility saliently resembles a third possibility: one in which Nogot drives a Ford he owns while Havit neither drives nor owns a car. This third possibility may not properly be ignored, because of the degree to which it is believed. This time, the resemblance is perfect so far as Havit is concerned, rather good so far as Nogot is concerned.

(2) The stopped clock is right twice a day. It says 4:39, as it has done for weeks. I look at it at 4:39; by luck I pick up a true belief. I have ignored the uneliminated possibility that I looked at it at 4:22 while it was stopped saying 4:39. That possibility was not properly ignored. It resembles actuality perfectly so far as the stopped clock goes.

(3) Unbeknownst to me, I am travelling in the land of the bogus barns; but my eye falls on one of the few real ones. I don't know that I am seeing a barn, because I may not properly ignore the possibility that I am seeing yet another of the abundant bogus barns. This possibility saliently resembles actuality in respect of the abundance of bogus barns, and the scarcity of real ones, hereabouts.

(4) Donald is in San Francisco, just as I have every reason to think he is. But, bent on deception, he is writing me letters and having them posted to me by his accomplice in Italy. If I had seen the phoney

knowledge, it is not normally counted among the Gettier problems. It is interesting to find that it yields to the same remedy.

letters, with their Italian stamps and postmarks, I would have concluded that Donald was in Italy. Luckily, I have not yet seen any of them. I ignore the uneliminated possibility that Donald has gone to Italy and is sending me letters from there. But this possibility is not properly ignored, because it resembles actuality both with respect to the fact that the letters are coming to me from Italy and with respect to the fact that those letters come, ultimately, from Donald. So I don't know that Donald is in San Francisco.

Next, there is the *Rule of Reliability*. This time, we have a presumptive rule about what *may* be properly ignored; and it is by means of this rule that we capture what is right about causal or reliabilist theories of knowing. Consider processes whereby information is transmitted to us: perception, memory, and testimony. These processes are fairly reliable.¹⁵ Within limits, we are entitled to take them for granted. We may properly presuppose that they work without a glitch in the case under consideration. Defeasibly – *very* defeasibly! – a possibility in which they fail may properly be ignored.

My visual experience, for instance, depends causally on the scene before my eyes, and what I believe about the scene before my eyes depends in turn on my visual experience. Each dependence covers a wide and varied range of alternatives.¹⁶ Of course, it is possible to hallucinate – even to hallucinate in such a way that all my perceptual experience and memory would be just as they actually are. That possibility never can be eliminated. But it can be ignored. And if it is properly ignored – as it mostly is – then vision gives me knowledge. Sometimes, though, the possibility of hallucination is not properly ignored; for sometimes we really do hallucinate. The Rule of Reliability may be defeated by the Rule of Actuality. Or it may be defeated by the Rules of Actuality and of Resemblance working together, in a

15 See Alvin Goldman, 'A Causal Theory of Knowing', *The Journal of Philosophy* 64 (1967) pp. 357–372; D. M. Armstrong, *Belief, Truth and Knowledge* (Cambridge: Cambridge University Press, 1973).

16 See my 'Veridical Hallucination and Prosthetic Vision', *Australasian Journal of Philosophy* 58 (1980) pp. 239–249. John Bigelow has proposed to model knowledge-delivering processes generally on those found in vision.

Gettier problem: if I am not hallucinating, but unbeknownst to me I live in a world where people mostly do hallucinate and I myself have only narrowly escaped, then the uneliminated possibility of hallucination is too close to actuality to be properly ignored.

We do not, of course, presuppose that nowhere ever is there a failure of, say, vision. The general presupposition that vision is reliable consists, rather, of a standing disposition to presuppose, concerning whatever particular case may be under consideration, that we have no failure in that case.

In similar fashion, we have two permissive *Rules of Method*. We are entitled to presuppose – again, very defeasibly – that a sample is representative; and that the best explanation of our evidence is the true explanation. That is, we are entitled properly to ignore possible failures in these two standard methods of non-deductive inference. Again, the general rule consists of a standing disposition to presuppose reliability in whatever particular case may come before us.

Yet another permissive rule is the *Rule of Conservatism*. Suppose that those around us normally do ignore certain possibilities, and it is common knowledge that they do. (They do, they expect each other to, they expect each other to expect each other to, . . .) Then – again, very defeasibly! – these generally ignored possibilities may properly be ignored. We are permitted, defeasibly, to adopt the usual and mutually expected presuppositions of those around us.

(It is unclear whether we need all four of these permissive rules. Some might be subsumed under others. Perhaps our habits of treating samples as representative, and of inferring to the best explanation, might count as normally reliable processes of transmission of information. Or perhaps we might subsume the Rule of Reliability under the Rule of Conservatism, on the ground that the reliable processes whereby we gain knowledge are familiar, are generally relied upon, and so are generally presupposed to be normally reliable. Then the only extra work done by the Rule of Reliability would be to cover less familiar – and merely hypothetical? – reliable processes, such as processes that relied on extrasensory faculties. Likewise, *mutatis mutandis*, we might subsume the Rules of Method under the Rule of Conservatism. Or we might instead think to subsume the Rule of Conservatism

under the Rule of Reliability, on the ground that what is generally presupposed tends for the most part to be true, and the reliable processes whereby this is so are covered already by the Rule of Reliability. Better redundancy than incompleteness, though. So, leaving the question of redundancy open, I list all four rules.)

Our final rule is the *Rule of Attention*. But it is more a triviality than a rule. When we say that a possibility is properly ignored, we mean exactly that; we do not mean that it *could have been* properly ignored. Accordingly, a possibility not ignored at all is *ipso facto* not properly ignored. What is and what is not being ignored is a feature of the particular conversational context. No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in *this* context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative. It is in the contextually determined domain. If it is an uneliminated possibility in which not-*P*, then it will do as a counter-example to the claim that *P* holds in every possibility left uneliminated by *S*'s evidence. That is, it will do as a counter-example to the claim that *S* knows that *P*.

Do some epistemology. Let your fantasies rip. Find uneliminated possibilities of error everywhere. Now that you are attending to them, just as I told you to, you are no longer ignoring them, properly or otherwise. So you have landed in a context with an enormously rich domain of potential counter-examples to ascriptions of knowledge. In such an extraordinary context, with such a rich domain, it never can happen (well, hardly ever) that an ascription of knowledge is true. Not an ascription of knowledge to yourself (either to your present self or to your earlier self, untainted by epistemology); and not an ascription of knowledge to others. That is how epistemology destroys knowledge. But it does so only temporarily. The pastime of epistemology does not plunge us forevermore into its special context. We can still do a lot of proper ignoring, a lot of knowing, and a lot of true ascribing of knowledge to ourselves and others, the rest of the time.

What is epistemology all about? The epistemology we've just been doing, at any rate, soon became an investigation of the ignoring of possibilities. But to investigate the ignoring of them was *ipso facto* not to ignore them. Unless this investigation of ours was an altogether

atypical sample of epistemology, it will be inevitable that epistemology must destroy knowledge. That is how knowledge is elusive. Examine it, and straightway it vanishes.

Is resistance useless? If you bring some hitherto ignored possibility to our attention, then straightway we are not ignoring it at all, so *a fortiori* we are not properly ignoring it. How can this alteration of our conversational state be undone? If you are persistent, perhaps it cannot be undone – at least not so long as you are around. Even if we go off and play backgammon, and afterward start our conversation afresh, you might turn up and call our attention to it all over again.

But maybe you called attention to the hitherto ignored possibility by mistake. You only suggested that we ought to suspect the butler because you mistakenly thought him to have a criminal record. Now that you know he does not – that was the *previous* butler – you wish you had not mentioned him at all. You know as well as we do that continued attention to the possibility you brought up impedes our shared conversational purposes. Indeed, it may be common knowledge between you and us that we would all prefer it if this possibility could be dismissed from our attention. In that case we might quickly strike a tacit agreement to speak just as if we were ignoring it; and after just a little of that, doubtless it really would be ignored.

Sometimes our conversational purposes are not altogether shared, and it is a matter of conflict whether attention to some far-fetched possibility would advance them or impede them. What if some far-fetched possibility is called to our attention not by a sceptical philosopher, but by counsel for the defence? We of the jury may wish to ignore it, and wish it had not been mentioned. If we ignored it now, we would bend the rules of cooperative conversation; but we may have good reason to do exactly that. (After all, what matters most to us as jurors is not whether we can truly be said to know; what really matters is what we should believe to what degree, and whether or not we should vote to convict.) We would ignore the far-fetched possibility if we could – but can we? Perhaps at first our attempted ignoring would be make-believe ignoring, or self-deceptive ignoring; later, perhaps, it might ripen into genuine ignoring. But in the meantime, do we know? There may be no definite answer. We are bending the rules,

and our practices of context-dependent attributions of knowledge were made for contexts with the rules unbent.

If you are still a contented fallibilist, despite my plea to hear the sceptical argument afresh, you will probably be discontented with the Rule of Attention. You will begrudge the sceptic even his very temporary victory. You will claim the right to resist his argument not only in everyday contexts, but even in those peculiar contexts in which he (or some other epistemologist) busily calls your attention to far-fetched possibilities of error. Further, you will claim the right to resist without having to bend any rules of cooperative conversation. I said that the Rule of Attention was a triviality: that which is not ignored at all is not properly ignored. But the Rule was trivial only because of how I had already chosen to state the *sotto voce* proviso. So you, the contented fallibilist, will think it ought to have been stated differently. Thus, perhaps: ‘Psst! – except for those possibilities we *could* properly have ignored’. And then you will insist that those far-fetched possibilities of error that we attend to at the behest of the sceptic are nevertheless possibilities we *could* properly have ignored. You will say that no amount of attention can, by itself, turn them into relevant alternatives.

If you say this, we have reached a standoff. I started with a puzzle: how can it be, when his conclusion is so silly, that the sceptic’s argument is so irresistible? My Rule of Attention, and the version of the proviso that made that Rule trivial, were built to explain how the sceptic manages to sway us – why his argument seems irresistible, however temporarily. If you continue to find it eminently resistible in all contexts, you have no need of any such explanation. We just disagree about the explanandum phenomenon.

I say *S* knows that *P* iff *P* holds in every possibility left uneliminated by *S*’s evidence – Psst! – except for those possibilities that *we* are properly ignoring. ‘We’ means: the speaker and hearers of a given context; that is, those of us who are discussing *S*’s knowledge together. It is our ignorings, not *S*’s own ignorings, that matter to what we can truly say about *S*’s knowledge. When we are talking about our own knowledge or ignorance, as epistemologists so often do, this is a distinction without a difference. But what if we are talking about someone else?

Suppose we are detectives; the crucial question for our solution of the crime is whether *S* already *knew*, when he bought the gun, that he was vulnerable to blackmail. We conclude that he did. *We* ignore various far-fetched possibilities, as hard-headed detectives should. But *S* does not ignore them. *S* is by profession a sceptical epistemologist. He never ignores much of anything. If it is our own ignorings that matter to the truth of our conclusion, we may well be right that *S* already knew. But if it is *S*'s ignorings that matter, then we are wrong, because *S* never knew much of anything. I say we may well be right; so it is our own ignorings that matter, not *S*'s.

But suppose instead that we are epistemologists considering what *S* knows. If we are well-informed about *S* (or if we are considering a well-enough specified hypothetical case), then if *S* attends to a certain possibility, we attend to *S*'s attending to it. But to attend to *S*'s attending to it is *ipso facto* to attend to it ourselves. In that case, unlike the case of the detectives, the possibilities we are properly ignoring must be among the possibilities that *S* himself ignores. We may ignore fewer possibilities than *S* does, but not more.

Even if *S* himself is neither sceptical nor an epistemologist, he may yet be clever at thinking up far-fetched possibilities that are uneliminated by his evidence. Then again, we well-informed epistemologists who ask what *S* knows will have to attend to the possibilities that *S* thinks up. Even if *S*'s idle cleverness does not lead *S* himself to draw sceptical conclusions, it nevertheless limits the knowledge that we can truly ascribe to him when attentive to his state of mind. More simply: his cleverness limits his knowledge. He would have known more, had he been less imaginative.¹⁷

Do I claim you can know *P* just by presupposing it?! Do I claim you can know that a possibility *W* does not obtain just by ignoring it? Is that not what my analysis implies, provided that the presupposing and

17 See Catherine Elgin, 'The Epistemic Efficacy of Stupidity', *Synthese* 74 (1988) pp. 297–311. The 'efficacy' takes many forms; some to do with knowledge (under various rival analyses), some to do with justified belief. See also Michael Williams, *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism* (Oxford: Blackwell, 1991) pp. 352–355, on the instability of knowledge under reflection.

the ignoring are proper? Well, yes. And yet I do not claim it. Or rather, I do not claim it for any specified *P* or *W*. I have to grant, in general, that knowledge just by presupposing and ignoring *is* knowledge; but it is an *especially* elusive sort of knowledge, and consequently it is an unclaimable sort of knowledge. You do not even have to practise epistemology to make it vanish. Simply *mentioning* any particular case of this knowledge, aloud or even in silent thought, is a way to attend to the hitherto ignored possibility, and thereby render it no longer ignored, and thereby create a context in which it is no longer true to ascribe the knowledge in question to yourself or others. So, just as we should think, presuppositions alone are not a basis on which to *claim* knowledge.

In general, when *S* knows that *P* some of the possibilities in which not-*P* are eliminated by *S*'s evidence and others of them are properly ignored. There are some that can be eliminated, but cannot properly be ignored. For instance, when I look around the study without seeing Possum the cat, I thereby eliminate various possibilities in which Possum is in the study; but had those possibilities not been eliminated, they could not properly have been ignored. And there are other possibilities that never can be eliminated, but can properly be ignored. For instance, the possibility that Possum is on the desk but has been made invisible by a deceiving demon falls normally into this class (though not when I attend to it in the special context of epistemology).

There is a third class: not-*P* possibilities that might either be eliminated or ignored. Take the far-fetched possibility that Possum has somehow managed to get into a closed drawer of the desk – maybe he jumped in when it was open, then I closed it without noticing him. That possibility could be eliminated by opening the drawer and making a thorough examination. But if uneliminated, it may nevertheless be ignored, and in many contexts that ignoring would be proper. If I look all around the study, but without checking the closed drawers of the desk, I may truly be said to know that Possum is not in the study – or at any rate, there are many contexts in which that may truly be said. But if I did check all the closed drawers, then I would know *better* that Possum is not in the study. My knowledge would be better in

the second case because it would rest more on the elimination of not-*P* possibilities, less on the ignoring of them.^{18,19}

Better knowledge is more stable knowledge: it stands more chance of surviving a shift of attention in which we begin to attend to some of the possibilities formerly ignored. If, in our new shifted context, we ask what knowledge we may truly ascribe to our earlier selves, we may find that only the better knowledge of our earlier selves still deserves the name. And yet, if our former ignorings were proper at the time, even the worse knowledge of our earlier selves could truly have been called knowledge in the former context.

Never – well, hardly ever – does our knowledge rest entirely on elimination and not at all on ignoring. So hardly ever is it quite as good as we might wish. To that extent, the lesson of scepticism is right – and right permanently, not just in the temporary and special context of epistemology.²⁰

What is it all for? Why have a notion of knowledge that works in the way I described? (Not a compulsory question. Enough to observe that

18 Mixed cases are possible: Fred properly ignores the possibility *W*₁ which Ted eliminates; however, Ted properly ignores the possibility *W*₂ which Fred eliminates. Ted has looked in all the desk drawers but not the file drawers, whereas Fred has checked the file drawers but not the desk. Fred's knowledge that Possum is not in the study is better in one way, Ted's is better in another.

19 To say truly that *X* is known, I must be properly ignoring any uneliminated possibilities in which not-*X*; whereas to say truly that *Y* is better known than *X*, I must be attending to some such possibilities. So I cannot say both in a single context. If I say 'X is known, but *Y* is better known', the context changes in mid-sentence: some previously ignored possibilities must stop being ignored. That can happen easily. Saying it the other way around – 'Y is better known than *X*, but even *X* is known' – is harder, because we must suddenly start to ignore previously unignored possibilities. That cannot be done, really; but we could bend the rules and make believe we had done it, and no doubt we would be understood well enough. Saying 'X is flat, but *Y* is flatter' (that is, 'X has no bumps at all, but *Y* has even fewer or smaller bumps') is a parallel case. And again, 'Y is flatter, but even *X* is flat' sounds clearly worse – but not altogether hopeless.

20 Thanks here to Stephen Hetherington. While his own views about better and worse knowledge are situated within an analysis of knowledge quite unlike mine, they withstand transplantation.

we do have it.) But I venture the guess that it is one of the messy short-cuts – like satisficing, like having indeterminate degrees of belief – that we resort to because we are not smart enough to live up to really high, perfectly Bayesian, standards of rationality. You cannot maintain a record of exactly which possibilities you have eliminated so far, much as you might like to. It is easier to keep track of which possibilities you have eliminated if you – Psst! – ignore many of all the possibilities there are. And besides, it is easier to list some of the propositions that are true in *all* the uneliminated, unignored possibilities than it is to find propositions that are true in *all and only* the uneliminated, unignored possibilities.

If you doubt that the word ‘know’ bears any real load in science or in metaphysics, I partly agree. The serious business of science has to do not with knowledge *per se*; but rather, with the elimination of possibilities through the evidence of perception, memory, etc., and with the changes that one’s belief system would (or might or should) undergo under the impact of such eliminations. Ascriptions of knowledge to yourself or others are a very sloppy way of conveying very incomplete information about the elimination of possibilities. It is as if you had said:

The possibilities eliminated, whatever else they may also include, at least include all the not-*P* possibilities; or anyway, all of those except for some we are presumably prepared to ignore just at the moment.

The only excuse for giving information about what really matters in such a sloppy way is that at least it is easy and quick! But it *is* easy and quick; whereas giving full and precise information about which possibilities have been eliminated seems to be extremely difficult, as witness the futile search for a ‘pure observation language’. If I am right about how ascriptions of knowledge work, they are a handy but humble approximation. They may yet be indispensable in practice, in the same way that other handy and humble approximations are.

If we analyse knowledge as a modality, as we have done, we cannot escape the conclusion that knowledge is closed under (strict) impli-

cation.²¹ Dretske has denied that knowledge is closed under implication; further, he has diagnosed closure as the fallacy that drives arguments for scepticism. As follows: the proposition that I have hands implies that I am not a handless being, and a *fortiori* that I am not a handless being deceived by a demon into thinking that I have hands. So, by the closure principle, the proposition that I know I have hands implies that I know that I am not handless and deceived. But I don't know that I am not handless and deceived – for how can I eliminate that possibility? So, by *modus tollens*, I don't know that I have hands. Dretske's advice is to resist scepticism by denying closure. He says that although having hands *does* imply not being handless and deceived, yet knowing that I have hands *does not* imply knowing that I am not handless and deceived. I do know the former, I do not know the latter.²²

What Dretske says is close to right, but not quite. Knowledge *is* closed under implication. Knowing that I have hands *does* imply knowing that I am not handless and deceived. Implication preserves truth – that is, it preserves truth in any given, fixed context. But if we switch contexts midway, all bets are off. I say (1) pigs fly; (2) what I just said had fewer than three syllables (true); (3) what I just said had fewer than four syllables (false). So 'less than three' does not imply 'less than four'? No! The context switched midway, the semantic value of the context-dependent phrase 'what I just said' switched with it. Likewise in the sceptical argument the context switched midway, and the semantic value of the context-dependent word 'know' switched with it. The premise 'I know that I have hands' was true in its everyday context, where the possibility of deceiving demons was properly ignored. The mention of that very possibility switched the context midway. The

21 A proof-theoretic version of this closure principle is common to all 'normal' modal logics: if the logic validates an inference from zero or more premises to a conclusion, then also it validates the inference obtained by prefixing the necessity operator to each premise and to the conclusion. Further, this rule is all we need to take us from classical sentential logic to the least normal modal logic. See Brian Chellas, *Modal Logic: An Introduction* (Cambridge: Cambridge University Press, 1980) p. 114.

22 Dretske, 'Epistemic Operators'. My reply follows the lead of Stine, 'Skepticism, Relevant Alternatives, and Deductive Closure', *op. cit.*; and (more closely) Cohen, 'How to be a Fallibilist', *op. cit.*

conclusion ‘I know that I am not handless and deceived’ was false in its context, because that was a context in which the possibility of deceiving demons was being mentioned, hence was not being ignored, hence was not being properly ignored. Dretske gets the phenomenon right, and I think he gets the diagnosis of scepticism right; it is just that he misclassifies what he sees. He thinks it is a phenomenon of logic, when really it is a phenomenon of pragmatics. Closure, rightly understood, survives the test. If we evaluate the conclusion for truth not with respect to the context in which it was uttered, but instead with respect to the different context in which the premise was uttered, then truth is preserved. And if, *per impossibile*, the conclusion could have been said in the same unchanged context as the premise, truth would have been preserved.

A problem due to Saul Kripke turns upon the closure of knowledge under implication. P implies that any evidence against P is misleading. So, by closure, whenever you know that P , you know that any evidence against P is misleading. And if you know that evidence is misleading, you should pay it no heed. Whenever we know – and we know a lot, remember – we should not heed any evidence tending to suggest that we are wrong. But that is absurd. Shall we dodge the conclusion by denying closure? I think not. Again, I diagnose a change of context. At first, it was stipulated that S knew, whence it followed that S was properly ignoring all possibilities of error. But as the story continues, it turns out that there is evidence on offer that points to some particular possibility of error. Then, by the Rule of Attention, that possibility is no longer properly ignored, either by S himself or by we who are telling the story of S . The advent of that evidence destroys S ’s knowledge, and thereby destroys S ’s licence to ignore the evidence lest he be misled.

There is another reason, different from Dretske’s, why we might doubt closure. Suppose two or more premises jointly imply a conclusion. Might not someone who is compartmentalized in his thinking – as we all are – know each of the premises but fail to bring them together in a single compartment? Then might he not fail to know the conclusion? Yes; and I would not like to plead idealization-of-rationality as an excuse for ignoring such cases. But I suggest that we might take not the whole compartmentalized thinker, but rather each

of his several overlapping compartments, as our ‘subjects’. That would be the obvious remedy if his compartmentalization amounted to a case of multiple personality disorder; but maybe it is right for milder cases as well.²³

A compartmentalized thinker who indulges in epistemology can destroy his knowledge, yet retain it as well. Imagine two epistemologists on a bushwalk. As they walk, they talk. They mention all manner of far-fetched possibilities of error. By attending to these normally ignored possibilities they destroy the knowledge they normally possess. Yet all the while they know where they are and where they are going! How so? The compartment in charge of philosophical talk attends to far-fetched possibilities of error. The compartment in charge of navigation does not. One compartment loses its knowledge, the other retains its knowledge. And what does the entire compartmentalized thinker know? Not an altogether felicitous question. But if we need an answer, I suppose the best thing to say is that *S* knows that *P* iff any one of *S*'s compartments knows that *P*. Then we can say what we would offhand want to say: yes, our philosophical bushwalkers still know their whereabouts.

Context-dependence is not limited to the ignoring and non-ignoring of far-fetched possibilities. Here is another case. Pity poor Bill! He squanders all his spare cash on the pokies, the races, and the lottery. He will be a wage slave all his days. We know he will never be rich. But if he wins the lottery (if he wins big), then he will be rich. Contrapositively: his never being rich, plus other things we know, imply that he will lose. So, by closure, if we know that he will never be rich, we know that he will lose. But when we discussed the case before, we concluded that we cannot know that he will lose. All the possibilities in which Bill loses and someone else wins saliently resemble the possibility in which Bill wins and the others lose; one of those possibilities is actual; so by the Rules of Actuality and of Resemblance, we may not properly ignore the possibility that Bill wins. But there is a loop-hole: the resemblance was required to be salient. Salience, as well as ignoring, may vary between contexts. Before, when I was explaining

23 See Stalnaker, *Inquiry*, pp. 79–99.

how the Rule of Resemblance applied to lotteries, I saw to it that the resemblance between the many possibilities associated with the many tickets was sufficiently salient. But this time, when we were busy pitting poor Bill for his habits and not for his luck, the resemblance of the many possibilities was not so salient. At that point, the possibility of Bill's winning was properly ignored; so then it was true to say that we knew he would never be rich. Afterward I switched the context. I mentioned the possibility that Bill might win, wherefore that possibility was no longer properly ignored. (Maybe there were two separate reasons why it was no longer properly ignored, because maybe I also made the resemblance between the many possibilities more salient.) It was true at first that we knew that Bill would never be rich. And at that point it was also true that we knew he would lose – but that was only true so long as it remained unsaid! (And maybe unthought as well.) Later, after the change in context, it was no longer true that we knew he would lose. At that point, it was also no longer true that we knew he would never be rich.

But wait. Don't you smell a rat? Haven't I, by my own lights, been saying what cannot be said? (Or whistled either.) If the story I told was true, how have I managed to tell it? In trendyspeak, is there not a problem of reflexivity? Does not my story deconstruct itself?

I said: S knows that P iff S 's evidence eliminates every possibility in which $\text{not-}P$ – Psst! – except for those possibilities that we are properly ignoring. That 'psst' marks an attempt to do the impossible – to mention that which remains unmentioned. I am sure you managed to make believe that I had succeeded. But I could not have done.

And I said that when we do epistemology, and we attend to the proper ignoring of possibilities, we make knowledge vanish. First we do know, then we do not. But I had been doing epistemology when I said that. The uneliminated possibilities were *not* being ignored – not just then. So by what right did I say even that we used to know?²⁴

²⁴ Worse still: by what right can I even say that we used to be in a position to say truly that we knew? Then, we were in a context where we properly ignored certain uneliminated possibilities of error. Now, we are in a context where we no longer ignore them. If *now* I comment retrospectively upon the truth of what was said

In trying to thread a course between the rock of fallibilism and the whirlpool of scepticism, it may well seem as if I have fallen victim to both at once. For do I not say that there are all those uneliminated possibilities of error? Yet do I not claim that we know a lot? Yet do I not claim that knowledge is, by definition, infallible knowledge?

I did claim all three things. But not all at once! Or if I did claim them all at once, that was an expository shortcut, to be taken with a pinch of salt. To get my message across, I bent the rules. If I tried to whistle what cannot be said, what of it? I relied on the cardinal principle of pragmatics, which overrides every one of the rules I mentioned: interpret the message to make it make sense – to make it consistent, and sensible to say.

When you have context-dependence, ineffability can be trite and unmysterious. Hush! [moment of silence] I might have liked to say, just then, ‘All of us are silent’. It was true. But I could not have said it truly, or whistled it either. For by saying it aloud, or by whistling, I would have rendered it false.

I could have said my say fair and square, bending no rules. It would have been tiresome, but it could have been done. The secret would have been to resort to ‘semantic ascent’. I could have taken great care to distinguish between (1) the language I use when I talk about knowledge, or whatever, and (2) the second language that I use to talk about the semantic and pragmatic workings of the first language. If you want to hear my story told that way, you probably know enough to do the job for yourself. If you can, then my informal presentation has been good enough.

then, which context governs: the context now or the context then? I doubt there is any general answer, apart from the usual principle that we should interpret what is said so as to make the message make sense.

